

Estimating Dynamic Treatment Regimes in Mobile Health Using V-learning

Daniel J. Lockett¹, Eric B. Laber², Anna R. Kahkoska³,
David M. Maahs⁴, Elizabeth Mayer-Davis³, Michael R. Kosorok¹

¹ Department of Biostatistics, University of North Carolina, Chapel Hill, NC 27599

² Department of Statistics, North Carolina State University, Raleigh, NC 27695

³ Department of Nutrition, University of North Carolina, Chapel Hill, NC 27599

⁴ Department of Pediatrics, Stanford University, Stanford, CA 94305

Abstract

The vision for precision medicine is to use individual patient characteristics to inform a personalized treatment plan that leads to the best healthcare possible for each patient. Mobile technologies have an important role to play in this vision as they offer a means to monitor a patient's health status in real-time and subsequently to deliver interventions if, when, and in the dose that they are needed. Dynamic treatment regimes formalize individualized treatment plans as sequences of decision rules, one per stage of clinical intervention, that map current patient information to a recommended treatment. However, existing methods for estimating optimal dynamic treatment regimes are designed for a small number of fixed decision points occurring on a coarse time-scale. We propose a new reinforcement learning method for estimating an optimal treatment regime that is applicable to data collected using mobile technologies in an outpatient setting. The proposed method accommodates an indefinite time horizon and minute-by-minute decision making that are common in mobile health applications. We show the proposed estimators are consistent and asymptotically normal under mild conditions. The proposed methods are applied to estimate an optimal dynamic treatment regime for controlling blood glucose levels in patients with type 1 diabetes.

1 Introduction

The use of mobile devices in clinical care, called mobile health (mHealth), provides an effective and scalable platform to assist patients in managing their illness (Free et al., 2013; Steinhubl et al., 2013). Advantages of mHealth interventions include real-time communication between a patient and their health-care provider as well as systems for delivering training, teaching, and social support (Kumar et al., 2013). Mobile technology can also be used to collect rich longitudinal data to estimate optimal dynamic treatment regimes and to deliver treatment that is deeply tailored to the individual. We propose a new reinforcement learning method to estimate dynamic treatment regimes in mHealth applications.

A dynamic treatment regime provides a framework to administer individualized treatment over time through a series of decision rules. Dynamic treatment regimes have been well-studied in the statistical and biomedical literature (Murphy, 2003; Robins, 2004; Moodie et al., 2007) and furthermore, statistical considerations in mHealth have been studied by, for example, Liao et al. (2015) and Klasnja et al. (2015). Although mobile technology has been successfully utilized in clinical areas such as diabetes (Quinn et al., 2011; Maahs et al., 2012), smoking cessation (Ali et al., 2012), and obesity (Bexelius et al., 2010), mHealth poses some unique challenges that preclude direct application of existing methodologies for dynamic treatment regimes. For example, mHealth applications typically involve a large number of time points per individual and no definite time horizon; the momentary signal may be weak and may not directly measure the outcome of interest; and estimation of optimal treatment strategies must be done online as data accumulate.

This work is motivated in part by our involvement in a study of mHealth as a management tool for type 1 diabetes. Type 1 diabetes is an autoimmune disease wherein the pancreas produces insufficient levels of insulin, a hormone needed to regulate blood glucose concentration. Patients with type 1 diabetes are forced to engage with management strategies including monitoring glucose levels, timing and dosing insulin injections, and regulating diet and physical activity. Increased glucose monitoring and attention to self management facilitate more frequent treatment adjustments and have been shown to improve patient outcomes (Levine et al., 2001; Haller et al., 2004; Ziegler et al., 2011). Thus, it is reasonable to believe that patient outcomes might be improved by diabetes management tools that are both deeply tailored to the individual and dynamic on a minute-by-minute basis to account for within patient variability over time. Mobile technology can be used to collect data on physical activity, glucose, and insulin at a fine granularity in an outpatient setting (Maahs et al., 2012). There is great potential for using this data to create comprehensive and accessible mHealth interventions for clinical use. We envision application of this work for use before the artificial pancreas (Weinzimer et al., 2008; Kowalski, 2015; Bergenstal et al., 2016) becomes widely available.

The sequential decision making process can be viewed as a Markov decision process (Puterman, 2014) and estimation of an optimal treatment regime can be accomplished with reinforcement learning algorithms such as Q-learning (Murphy, 2005; Zhao et al., 2009; Tang and Kosorok, 2012; Schulte et al., 2014; Ertefaie, 2014). We propose an alternative to Q-learning that is suited to mHealth applications. Our approach, which we call V-learning,

involves estimating the optimal policy among a pre-specified class of policies. It requires minimal assumptions about the data-generating model, allows for an indefinite time horizon, and permits estimating a randomized decision rule that can be implemented online.

In Section 2, we describe the setup and present our estimator for application to offline observational data. In Section 3, we extend our estimator for application to online estimation with accumulating data. Theoretical results, including consistency and asymptotic normality of the proposed estimators, are presented in Section 4. We apply the proposed method to the motivating example of type 1 diabetes management using simulated data in Section 5. A case study using data on patients with type 1 diabetes is presented in Section 6 and we conclude with a discussion in Section 7. Proofs of technical results are in the appendix.

2 Offline estimation from observational data

We assume that the available data are $\{(\mathbf{S}_i^1, A_i^1, \mathbf{S}_i^2, \dots, \mathbf{S}_i^{T_i}, A_i^{T_i}, \mathbf{S}_i^{T_i+1})\}_{i=1}^n$, which comprise n independent, identically distributed trajectories $(\mathbf{S}^1, A^1, \mathbf{S}^2, \dots, \mathbf{S}^T, A^T, \mathbf{S}^{T+1})$, where: $\mathbf{S}^t \in \mathbb{R}^p$ denotes a summary of patient information collected up to and including time t ; $A^t \in \mathcal{A}$ denotes the treatment assigned at time t ; and $T \in \mathbb{Z}_+$ denotes the (possibly random) patient follow-up time. We assume that the data generating model is a time-homogeneous Markov process so that $\mathbf{S}^{t+1} \perp\!\!\!\perp (A^{t-1}, \mathbf{S}^{t-1}, \dots, A^1, \mathbf{S}^1) | (A^t, \mathbf{S}^t)$ and the conditional density $p(\mathbf{S}^{t+1} | A^t, \mathbf{S}^t)$ is the same for all $t \geq 1$. Let $L^t \in \{0, 1\}$ denote an indicator that the patient is still in follow-up at time t , i.e., $L^t = 1$ if the patient is being followed at time t and zero otherwise. We assume that L^t is contained in \mathbf{S}^t so that $P(L^{t+1} = 1 | A^t, \mathbf{S}^t, \dots, A^1, \mathbf{S}^1) = P(L^{t+1} = 1 | A^t, \mathbf{S}^t)$ and $L^t = 0$ implies $L^{t+1} = 0$ with probability one. Furthermore, we assume a known utility function $u : \mathbb{R}^p \times \mathcal{A} \times \mathbb{R}^p \rightarrow \mathbb{R}$ so that $U^t = u(\mathbf{S}^{t+1}, A^t, \mathbf{S}^t)$ measures the ‘goodness’ of choosing treatment A^t in state \mathbf{S}^t and subsequently transitioning to state \mathbf{S}^{t+1} . The goal is to select treatments to maximize expected cumulative utility; treatment selection is formalized using a treatment regime (Schulte et al., 2014; Kosorok and Moodie, 2015) and the utility associated with any regime is defined using potential outcomes (Rubin, 1978).

Let $\mathcal{B}(\mathcal{A})$ denote the space of probability distributions over \mathcal{A} . A treatment regime in this context is a function $\pi : \text{dom } \mathbf{S}^t \rightarrow \mathcal{B}(\mathcal{A})$ so that, under π , a decision maker presented with state $\mathbf{S}^t = \mathbf{s}^t$ at time t will select action $a^t \in \mathcal{A}$ with probability $\pi(a^t; \mathbf{s}^t)$. Define $\bar{a}^t = (a^1, \dots, a^t) \in \mathcal{A}^t$, and $\bar{a}^\infty = (a^1, a^2, \dots) \in \mathcal{A}^\infty$. The set of potential outcomes is

$$\mathbf{W}^* = \left\{ \mathbf{S}^1, \mathbf{S}^{*2}(a^1), \dots, \mathbf{S}^{*T^*(\bar{a}^\infty)}(\bar{a}^{T^*(\bar{a}^\infty)-1}) : \right. \\ \left. T^*(\bar{a}^\infty) = \inf \{t \geq 1 : L^{*t}(\bar{a}^{t-1}) = 0\}, \bar{a}^\infty \in \mathcal{A}^\infty \right\},$$

where $\mathbf{S}^{*t}(\bar{a}^{t-1})$ is the potential state and $L^{*t}(\bar{a}^{t-1})$ is the potential follow-up status at time t under treatment sequence \bar{a}^{t-1} . Thus, the potential utility at time t is $U^{*t}(\bar{a}^t) = u\left\{\mathbf{S}^{*(t+1)}(\bar{a}^t), a^t, \mathbf{S}^{*t}(\bar{a}^{t-1})\right\}$. For any π , define $\{\xi_\pi^t(\cdot)\}_{t \geq 1}$ to be a sequence of independent,

\mathcal{A} -valued stochastic processes indexed by $\text{dom } \mathbf{S}^t$ such that $P\{\xi_\pi^t(\mathbf{s}^t) = a^t\} = \pi(a^t; \mathbf{s}^t)$. The potential follow-up time under π is

$$T^*(\pi) = \sum_{t \geq 1} \sum_{\bar{\mathbf{a}}^t \in \mathcal{A}^t} t 1\{\sup_{\bar{\mathbf{a}}^{t+1}} T^*(\bar{\mathbf{a}}^t, \bar{\mathbf{a}}^{t+1}) = t\} \prod_{v=1}^t 1[\xi_\pi^v\{\mathbf{S}^{*v}(\bar{\mathbf{a}}^{v-1})\} = a^v],$$

where $\bar{\mathbf{a}}^{t+1} = (a^{t+1}, a^{t+2}, \dots)$. The potential utility under π at time t is

$$U^{*t}(\pi) = \begin{cases} \sum_{\bar{\mathbf{a}}^t \in \mathcal{A}^t} U^{*t}(\bar{\mathbf{a}}^t) \prod_{v=1}^t 1[\xi_\pi^v\{\mathbf{S}^{*v}(\bar{\mathbf{a}}^{v-1})\} = a^v], & \text{if } T^*(\pi) \geq t \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathbf{S}^{*1}(\bar{\mathbf{a}}^0) = \mathbf{S}^1$. Thus, utility is set to zero after a patient is lost to follow-up. However, in certain situations, utility may be constructed so as to take a negative value at the time point when the patient is lost to follow-up, e.g., if the patient discontinues treatment because of a negative effect associated with the intervention. Define the state-value function $V(\pi, \mathbf{s}^t) = \mathbb{E}\{\sum_{k \geq 0} \gamma^k U^{*(t+k)}(\pi) | \mathbf{S}^t = \mathbf{s}^t\}$ (Sutton and Barto, 1998), where $\gamma \in (0, 1)$ is a fixed constant that captures the trade-off between short- and long-term outcomes. For any distribution \mathcal{R} on $\text{dom } \mathbf{S}^1$, define the value function with respect to reference distribution \mathcal{R} as $V_{\mathcal{R}}(\pi) = \int V(\pi, \mathbf{s}) d\mathcal{R}(\mathbf{s})$; throughout, we assume that this reference distribution is fixed. The reference distribution can be thought of as a distribution of initial states and we estimate it from the data in the implementation in Sections 5 and 6. For a pre-specified class of regimes, Π , the optimal regime, $\pi_{\mathcal{R}}^{\text{opt}} \in \Pi$, satisfies $V_{\mathcal{R}}(\pi_{\mathcal{R}}^{\text{opt}}) \geq V_{\mathcal{R}}(\pi)$ for all $\pi \in \Pi$.

To construct an estimator of $\pi_{\mathcal{R}}^{\text{opt}}$, we make a series of assumptions that connect the potential outcomes in \mathbf{W}^* with the data-generating model.

Assumption 1. Strong ignorability, $A^t \perp\!\!\!\perp \mathbf{W}^* | \mathbf{S}^t$ for all t .

Assumption 2. Consistency, $\mathbf{S}^t = \mathbf{S}^{*t}(\bar{\mathbf{A}}^{t-1})$ for all t and $T = T^*(\bar{\mathbf{A}}^\infty)$.

Assumption 3. Positivity, there exists $c_0 > 0$ so that $P(A^t = a^t | \mathbf{S}^t = \mathbf{s}^t) \geq c_0$ for all $a^t \in \mathcal{A}$, $\mathbf{s}^t \in \text{dom } \mathbf{S}^t$, and all t .

In addition, we implicitly assume that there is no interference among the experimental units. These assumptions are common in the context of estimating dynamic treatment regimes (Robins, 2004; Schulte et al., 2014). Assumptions 1 and 3 hold by construction in a micro-randomized trial (Klasnja et al., 2015; Liao et al., 2015).

Let $\mu^t(a^t; \mathbf{s}^t) = P(A^t = a^t | \mathbf{S}^t = \mathbf{s}^t)$ for each $t \geq 1$. In a micro-randomized trial, $\mu^t(a^t; \mathbf{s}^t)$ is a known randomization probability; in an observational study, it must be estimated from the data. The following lemma characterizes $V_{\mathcal{R}}(\pi)$ for any regime, π , in terms of the data-generating model. A proof is provided in the appendix.

Lemma 2.1. *Let π denote an arbitrary regime and $\gamma \in (0, 1)$ a discount factor. Then, under assumptions 1-3 and provided interchange of the sum and integration is justified, the state-value function of π at \mathbf{s}^t is*

$$V(\pi, \mathbf{s}^t) = \sum_{k \geq 0} \mathbb{E} \left[\gamma^k U^{t+k} \left\{ \prod_{v=0}^k \frac{\pi(A^{v+t}; \mathbf{S}^{v+t})}{\mu^{v+t}(A^{v+t}; \mathbf{S}^{v+t})} \right\} \middle| \mathbf{S}^t = \mathbf{s}^t \right]. \quad (1)$$

The preceding result will form the basis for an estimating equation for $V_{\mathcal{R}}(\pi)$. Write the right hand side of (1) as

$$\begin{aligned} V(\pi, \mathbf{S}^t) &= \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \left(U^t + \gamma \sum_{k \geq 0} \mathbb{E} \left[\gamma^k U^{t+k+1} \left\{ \prod_{v=0}^k \frac{\pi(A^{v+t+1}; \mathbf{S}^{v+t+1})}{\mu^{v+t+1}(A^{v+t+1}; \mathbf{S}^{v+t+1})} \right\} \middle| \mathbf{S}^{t+1} \right] \right) \middle| \mathbf{S}^t \right\} \\ &= \mathbb{E} \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \{ U^t + \gamma V(\pi, \mathbf{S}^{t+1}) \} \middle| \mathbf{S}^t \right], \end{aligned}$$

from which it follows that

$$0 = \mathbb{E} \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \{ U^t + \gamma V(\pi, \mathbf{S}^{t+1}) - V(\pi, \mathbf{S}^t) \} \middle| \mathbf{S}^t \right].$$

Subsequently, for any function ψ defined on $\text{dom } \mathbf{S}^t$, the state-value function satisfies

$$0 = \mathbb{E} \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \{ U^t + \gamma V(\pi, \mathbf{S}^{t+1}) - V(\pi, \mathbf{S}^t) \} \psi(\mathbf{S}^t) \right], \quad (2)$$

which is an importance-weighted variant of the well-known Bellman optimality equation (Sutton and Barto, 1998).

Let $V(\pi, \mathbf{s}; \theta^\pi)$ denote a model for $V(\pi, \mathbf{s})$ indexed by $\theta^\pi \in \Theta \subseteq \mathbb{R}^q$. We assume that the map $\theta^\pi \mapsto V(\pi, \mathbf{s}; \theta^\pi)$ is differentiable everywhere for each fixed \mathbf{s} and π . Let $\nabla_{\theta^\pi} V(\pi, \mathbf{s}; \theta^\pi)$ denote the gradient of $V(\pi, \mathbf{s}; \theta^\pi)$ and define

$$\Lambda_n(\pi, \theta^\pi) = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^{T_i} \frac{\pi(A_i^t; \mathbf{S}_i^t)}{\mu^t(A_i^t; \mathbf{S}_i^t)} \{ U_i^t + \gamma V(\pi, \mathbf{S}_i^{t+1}; \theta^\pi) - V(\pi, \mathbf{S}_i^t; \theta^\pi) \} \nabla_{\theta^\pi} V(\pi, \mathbf{S}_i^t; \theta^\pi). \quad (3)$$

Given a positive definite matrix $\Omega \in \mathbb{R}^{q \times q}$ and penalty function $\mathcal{P} : \mathbb{R}^q \rightarrow \mathbb{R}_+$ define $\hat{\theta}_n^\pi = \arg \min_{\theta^\pi \in \Theta} \{ \Lambda_n(\pi, \theta^\pi)^\top \Omega \Lambda_n(\pi, \theta^\pi) + \lambda_n \mathcal{P}(\theta^\pi) \}$, where λ_n is a tuning parameter. Subsequently, $V(\pi, \mathbf{s}; \hat{\theta}_n^\pi)$ is the estimated state-value function under π in state \mathbf{s} . Thus, given a reference distribution, \mathcal{R} , the estimated value of a regime, π , is $\hat{V}_{n, \mathcal{R}}(\pi) = \int V(\pi, \mathbf{s}; \hat{\theta}_n^\pi) d\mathcal{R}(\mathbf{s})$ and the estimated optimal regime is $\hat{\pi}_n = \arg \max_{\pi \in \Pi} \hat{V}_{n, \mathcal{R}}(\pi)$.

3 Online estimation from accumulating data

Suppose we have accumulating data $\{(\mathbf{S}_i^1, A_i^1, \mathbf{S}_i^2, \dots)\}_{i=1}^n$, where \mathbf{S}_i^t and A_i^t represent the state and action for patient $i = 1, \dots, n$ at time $t \geq 1$. At each time $t \geq 1$, we estimate an optimal policy in a class, Π , using data collected up to time t , take actions according to the estimated optimal policy, and estimate a new policy using the resulting states. Let $\hat{\pi}_n^t$ be the estimated policy at time $t \geq 1$, i.e., $\hat{\pi}_n^t$ is estimated after observing state \mathbf{S}^{t+1} and before taking action A^{t+1} . If Π is a class of randomized policies, we can select an action for

a patient presenting with $\mathbf{S}^{t+1} = \mathbf{s}^{t+1}$ according to $\hat{\pi}_n^t(\cdot; \mathbf{s}^{t+1})$, i.e., we draw A^{t+1} according to the distribution $P(A^{t+1} = a) = \hat{\pi}_n^t(a; \mathbf{s}^{t+1})$. We will use a class of randomized policies in the following. However, if a class of deterministic policies is of interest, we can inject some randomness into $\hat{\pi}_n^t$ to facilitate exploration (Sutton and Barto, 1998). One example is an ϵ -greedy strategy (see Sutton and Barto, 1998), which selects the estimated optimal action with probability $1 - \epsilon$ and otherwise sample equally from all other actions. Because an ϵ -greedy strategy can be used to introduce randomness into a deterministic policy, we can assume a class of randomized policies.

At each time $t \geq 1$, let $\hat{\theta}_{n,t}^\pi = \arg \min_{\theta^\pi \in \Theta} \{\Lambda_{n,t}(\pi, \theta^\pi)^\top \Omega \Lambda_{n,t}(\pi, \theta^\pi) + \lambda_n \mathcal{P}(\theta^\pi)\}$, where Ω and \mathcal{P} are as defined in Section 2 and

$$\Lambda_{n,t}(\pi, \theta^\pi) = \frac{1}{n} \sum_{i=1}^n \sum_{v=1}^t \frac{\pi(A_i^v; \mathbf{S}_i^v)}{\hat{\pi}_n^{v-1}(A_i^v; \mathbf{S}_i^v)} \{U_i^v + \gamma V(\pi, \mathbf{S}_i^{v+1}; \theta^\pi) - V(\pi, \mathbf{S}_i^v; \theta^\pi)\} \nabla_{\theta^\pi} V(\pi, \mathbf{S}_i^v; \theta^\pi). \quad (4)$$

$\hat{\pi}_n^0$ is some initial randomized policy. We note that estimating equation (4) is similar to (3), except that $\hat{\pi}_n^{v-1}$ replaces μ^v as the data-generating policy. Given an estimate of the value of π at time t , $\hat{V}_{n,\mathcal{R},t}(\pi) = \int V(\pi, \hat{\theta}_{n,t}^\pi) d\mathcal{R}(\mathbf{s})$, the optimal policy at time t is $\hat{\pi}_n^t = \arg \max_{\pi \in \Pi} \hat{V}_{n,\mathcal{R},t}(\pi)$. An alternative way to encourage exploration through the action space is to choose $\hat{\pi}_n^t = \arg \max_{\pi \in \Pi} \{\hat{V}_{n,\mathcal{R},t}(\pi) + \alpha^t \hat{\psi}^t(\pi)\}$ for some sequence $\alpha^t \geq 0$, where $\hat{\psi}^t(\pi)$ is a measure of uncertainty in $\hat{V}_{n,\mathcal{R},t}(\pi)$. An example of this is upper confidence bound sampling (UCB) (Lai and Robbins, 1985).

4 Theoretical results

In this section, we establish asymptotic properties of $\hat{\theta}_n^\pi$ and $\hat{\pi}_n$. Throughout, we assume assumptions 1-3 of Section 2. Let $\Phi = (\phi_1, \dots, \phi_q)^\top$ be a vector of pre-specified basis functions and let $\Phi(\mathbf{s}_i^t) = \{\phi_1(\mathbf{s}_i^t), \dots, \phi_q(\mathbf{s}_i^t)\}^\top$. Assume $V(\pi, \mathbf{s}_i^t; \theta^\pi) = \Phi(\mathbf{s}_i^t)^\top \theta^\pi$. Computational efficiency is gained from the linearity of $V(\pi, \mathbf{s}_i^t; \theta^\pi)$ in θ^π ; flexibility can be achieved through the choice of Φ . Under this working model,

$$\begin{aligned} \Lambda_n(\pi, \theta^\pi) &= \left[n^{-1} \sum_{i=1}^n \sum_{t=1}^{T_i} \frac{\pi(A_i^t; \mathbf{S}_i^t)}{\mu^t(A_i^t; \mathbf{S}_i^t)} \{ \gamma \Phi(\mathbf{S}_i^t) \Phi(\mathbf{S}_i^{t+1})^\top - \Phi(\mathbf{S}_i^t) \Phi(\mathbf{S}_i^t)^\top \} \right] \theta^\pi \\ &\quad + n^{-1} \sum_{i=1}^n \sum_{t=1}^{T_i} \left\{ \frac{\pi(A_i^t; \mathbf{S}_i^t)}{\mu^t(A_i^t; \mathbf{S}_i^t)} U_i^t \Phi(\mathbf{S}_i^t) \right\}. \end{aligned} \quad (5)$$

Let $\hat{\theta}_n^\pi = \arg \min_{\theta^\pi \in \Theta} \{\Lambda_n(\pi, \theta^\pi)^\top \Omega \Lambda_n(\pi, \theta^\pi) + \lambda_n \mathcal{P}(\theta^\pi)\}$ for a known positive definite matrix, Ω , a known penalty function, $\mathcal{P} : \mathbb{R}^q \rightarrow \mathbb{R}_+$, and a tuning parameter, λ_n . For fixed π , denote the true θ^π by θ_0^π , i.e., $V(\pi, \mathbf{s}) = \Phi(\mathbf{s})^\top \theta_0^\pi$. Let $\nu = \int \Phi(\mathbf{s}) d\mathcal{R}(\mathbf{s})$ so that $V_{\mathcal{R}}(\pi) = \nu^\top \theta_0^\pi$. Define $\hat{V}_{n,\hat{\mathcal{R}}}(\pi) = \{\mathbb{E}_n \Phi(\mathbf{S})\}^\top \hat{\theta}_n^\pi$, where \mathbb{E}_n denotes the empirical measure of the observed data.

Our main results are summarized in Theorems 4.2 and 4.3 below. Because each patient trajectory is a stationary Markov chain, we need to use asymptotic theory based on stationary processes; consequently, some of the required technical conditions are more difficult to verify than those for i.i.d. data. Define the bracketing integral for a class of functions, \mathcal{F} , by $J_{[]} \{\delta, \mathcal{F}, L_r(P)\} = \int_0^\delta \sqrt{\log N_{[]} \{\epsilon, \mathcal{F}, L_r(P)\}} d\epsilon$, where the bracketing number for \mathcal{F} , $N_{[]} \{\epsilon, \mathcal{F}, L_r(P)\}$, is the number of $L_r(P)$ ϵ -brackets needed such that each element of \mathcal{F} is contained in at least one bracket (see Chapter 2 of Kosorok, 2008). For any stationary sequence of possibly dependent random variables, $\{X^t\}_{t \geq 1}$, let \mathcal{M}_a^b be the σ -field generated by X^a, \dots, X^b and define $\zeta(k) = \mathbb{E} [\sup_{m \geq 1} \{|P(B|\mathcal{M}_1^m) - P(B)| : B \in \mathcal{M}_{m+k}^\infty\}]$. We say that the chain $\{X^t\}_{t \geq 1}$ is absolutely regular if $\zeta(k) \rightarrow 0$ as $k \rightarrow \infty$ (also called β -mixing in Chapter 11 of Kosorok, 2008). Let $\Pi = \{\pi_\beta : \beta \in \mathcal{B}\}$ be a parametric class of policies. We make the following assumptions.

Assumption 4. There exists a $2 < \rho < \infty$ such that

1. $\mathbb{E}|U^t|^{3\rho} < \infty$, $\mathbb{E}\|\Phi(\mathbf{S}^t)\|^{3\rho} < \infty$, and $\mathbb{E}\|\mathbf{S}^t\|^{3\rho} < \infty$.
2. The sequence $\{(\mathbf{S}^t, A^t)\}_{t \geq 1}$ is absolutely regular with $\sum_{k=1}^\infty k^{2/(\rho-2)} \zeta(k) < \infty$.
3. The bracketing integral of the class of policies, $J_{[]} \{\infty, \Pi, L_{3\rho}(P)\} < \infty$.

Assumption 5. There exists some $c_1 > 0$ such that

$$\inf_{\pi \in \Pi} c^\top \mathbb{E} \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \{ \Phi(\mathbf{S}^t) \Phi(\mathbf{S}^t)^\top - \gamma^2 \Phi(\mathbf{S}^{t+1}) \Phi(\mathbf{S}^{t+1})^\top \} \right] c \geq c_1 \|c\|^2$$

for all $c \in \mathbb{R}^q$.

Assumption 6. The map $\beta \mapsto V_{\mathcal{R}}(\pi_\beta)$ has a unique and well separated maximum over β in the interior of \mathcal{B} ; let β_0 denote the maximizer.

Assumption 7. The following condition holds: $\sup_{\|\beta_1 - \beta_2\| \leq \delta} \mathbb{E} \|\pi_{\beta_1}(A; \mathbf{S}) - \pi_{\beta_2}(A; \mathbf{S})\| \rightarrow 0$ as $\delta \downarrow 0$.

One class of policies that satisfies assumption 7 is as follows. Define the class $\Pi = \{\pi_\beta : \beta \in \mathcal{B} \subset \mathbb{R}^{p(J-1)}\}$ by

$$\pi_\beta(a; \mathbf{s}) = \frac{a_J + \sum_{k=1}^{J-1} a_k \exp(\mathbf{s}^\top \beta_k)}{1 + \sum_{k=1}^{J-1} \exp(\mathbf{s}^\top \beta_k)}$$

for $\beta = (\beta_1^\top, \dots, \beta_{J-1}^\top)^\top$, where $\beta_j = (\beta_{j1}, \dots, \beta_{jp})^\top$ is a vector of coefficients for treatment j and \mathcal{B} is a known compact set. We assume that treatment is coded as $a = (a_1, \dots, a_J)^\top$ where $a_j = 1$ and $a_k = 0$ for all $k \neq j$ whenever treatment j is assigned. This class of randomized policies will be used in the implementation in Sections 5 and 6.

Remark 4.1. Assumption 4 requires certain finite moments and that the dependence between observations on the same patient vanishes as observations become further apart. In Lemma 7.2 in the appendix, we verify part 3 of assumption 4 and assumption 7 for the class of policies introduced above. Assumption 5 is needed to show the existence of a unique θ_0^π uniformly over Π and assumption 6 requires that the true optimal decision in each state is unique (similar to assumption A.8 of Ertefaie, 2014). Assumption 7 requires smoothness on the class of policies.

The main results of this section are stated below. Theorem 4.2 states that there exists a unique solution to $0 = \mathbb{E}\Lambda_n(\pi, \theta^\pi)$ uniformly over Π and that the estimator $\hat{\theta}_n$ converges weakly to a mean zero Gaussian process in $\ell^\infty(\Pi)$.

Theorem 4.2. *Under the given assumptions, the following hold.*

1. For all $\pi \in \Pi$, there exists a $\theta_0^\pi \in \mathbb{R}^q$ such that $\mathbb{E}\Lambda_n(\pi, \theta^\pi)$ has a zero at $\theta^\pi = \theta_0^\pi$. Moreover, $\sup_{\pi \in \Pi} \|\theta_0^\pi\| < \infty$ and $\sup_{\|\beta_1 - \beta_2\| \leq \delta} \|\theta_0^{\pi_{\beta_1}} - \theta_0^{\pi_{\beta_2}}\| \rightarrow 0$ as $\delta \downarrow 0$.
2. Let $\mathbb{G}(\pi)$ be a tight, mean zero Gaussian process indexed by Π with covariance $\mathbb{E}\{\mathbb{G}(\pi_1)\mathbb{G}(\pi_2)\} = w_1(\pi_1)^{-1}w_0(\pi_1, \pi_2)w_1(\pi_2)^{-\top}$ where

$$w_0(\pi_1, \pi_2) = \mathbb{E} \left[\frac{\pi_1(A^t; \mathbf{S}^t)\pi_2(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)^2} \{U^t + \gamma\Phi(\mathbf{S}^{t+1})\theta_0^{\pi_1} - \Phi(\mathbf{S}^t)\theta_0^{\pi_1}\} \right. \\ \left. \{U^t + \gamma\Phi(\mathbf{S}^{t+1})\theta_0^{\pi_2} - \Phi(\mathbf{S}^t)\theta_0^{\pi_2}\} \Phi(\mathbf{S}^t)\Phi(\mathbf{S}^t)^\top \right]$$

and

$$w_1(\pi) = \mathbb{E} \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \Phi(\mathbf{S}^t) \{ \Phi(\mathbf{S}^t) - \gamma\Phi(\mathbf{S}^{t+1}) \} \right].$$

Then, $\sqrt{n}(\hat{\theta}_n^\pi - \theta_0^\pi) \rightsquigarrow \mathbb{G}(\pi)$ in $\ell^\infty(\Pi)$.

3. Let $\mathbb{G}(\pi)$ be as defined in part 2. Then, $\sqrt{n}\{\hat{V}_{n, \hat{\mathcal{R}}}(\pi) - V_{\mathcal{R}}(\pi)\} \rightsquigarrow \nu^\top \mathbb{G}(\pi)$ in $\ell^\infty(\Pi)$.

Theorem 4.3 below gives us that the estimated optimal policy converges in probability to the true optimal policy over Π and that the estimated value of the estimated optimal policy converges to the true value of the estimated optimal policy.

Theorem 4.3. *Under the given assumptions, the following hold.*

1. Let $\hat{\beta}_n = \arg \max_{\beta \in \mathcal{B}} \hat{V}_{n, \hat{\mathcal{R}}}(\pi_\beta)$ and $\beta_0 = \arg \max_{\beta \in \mathcal{B}} V_{\mathcal{R}}(\pi_\beta)$. Then, $\|\hat{\beta}_n - \beta_0\| \xrightarrow{P} 0$.
2. Let $\sigma_0^2 = \nu^\top w_1(\pi_{\beta_0})^{-1}w_0(\pi_{\beta_0}, \pi_{\beta_0})w_1(\pi_{\beta_0})^{-\top}\nu$. Then, $\sqrt{n}\{\hat{V}_{n, \hat{\mathcal{R}}}(\pi_{\hat{\beta}_n}) - V_{\mathcal{R}}(\pi_{\hat{\beta}_n})\} \rightsquigarrow N(0, \sigma_0^2)$.

3. A consistent estimator for σ_0^2 is

$$\hat{\sigma}_n^2 = \{\mathbb{E}_n \Phi(\mathbf{S}^t)\}^\top \hat{w}_1(\pi_{\hat{\beta}_n})^{-1} \hat{w}_0(\pi_{\hat{\beta}_n}, \pi_{\hat{\beta}_n}) \hat{w}_1(\pi_{\hat{\beta}_n})^{-\top} \{\mathbb{E}_n \Phi(\mathbf{S}^t)\},$$

where

$$\hat{w}_1(\pi) = \mathbb{E}_n \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \Phi(\mathbf{S}^t) \{\Phi(\mathbf{S}^t) - \gamma \Phi(\mathbf{S}^{t+1})\} \right]$$

and

$$\begin{aligned} \hat{w}_0(\pi_1, \pi_2) &= \mathbb{E}_n \left[\frac{\pi_1(A^t; \mathbf{S}^t) \pi_2(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)^2} \left\{ U^t + \gamma \Phi(\mathbf{S}^{t+1}) \hat{\theta}_n^{\pi_1} - \Phi(\mathbf{S}^t) \hat{\theta}_n^{\pi_1} \right\} \right. \\ &\quad \left. \left\{ U^t + \gamma \Phi(\mathbf{S}^{t+1}) \hat{\theta}_n^{\pi_2} - \Phi(\mathbf{S}^t) \hat{\theta}_n^{\pi_2} \right\} \Phi(\mathbf{S}^t) \Phi(\mathbf{S}^t)^\top \right]. \end{aligned}$$

where \mathbb{E}_n denotes the empirical measure of the observed data.

Proofs of the above results are in the appendix along with a result on bracketing entropy that is needed for the proof of Theorem 4.2 and a proof that the class of policies introduced above satisfies the necessary bracketing integral assumption.

5 Simulation experiments

In this section, we simulate a cohort of youths with type 1 diabetes similar to that of Maahs et al. (2012). Maahs et al. (2012) collected data in an outpatient setting using mobile devices for $N = 31$ patients followed for five days. Blood glucose level was tracked continuously using continuous glucose monitoring, physical activity was tracked using accelerometers, insulin injections were administered and logged by an insulin pump or injected and logged by the patient, and food intake was self-reported by each patient via telephone-based interviews.

5.1 Observational data

We divide each day of follow-up into 60 minute intervals. Thus, for one day of follow-up, we observe $T = 24$ time points per simulated patient. We define the state for patient i at time t to be the vector containing average blood glucose and total physical activity over the previous two time points and total food intake over the previous four time points. Our hypothetical mHealth intervention consists of three possible actions which can be taken at each time point: an insulin injection administered via insulin pump, a suggestion for food intake delivered via mobile app, or a suggestion for physical activity delivered via mobile app; available treatments consist of all combinations of these. The utility at time t is defined as a weighted sum of hypo- and hyperglycemic episodes in the 60 minutes preceding and following time t . Weights are defined in Table 1. Utility at each time point ranges from -6 to 0 with larger values being more preferable. Weights were chosen to reflect the relative

weight	blood glucose
-3	Gl \leq 70
-1	70 < Gl \leq 80
0	80 < Gl \leq 120
-1	120 < Gl \leq 150
-2	Gl > 150

Table 1: Weights for determining utility

clinical consequences of high and low blood glucose. For example, acute hypoglycemia, characterized by blood glucose levels below 70 mg/dL, is an emergency that can result in coma or death.

V-learning requires defining a class of policies over which to maximize estimated value. In our setting, there are eight possible treatments, a_1, \dots, a_8 , that can be administered at each time point, corresponding to all possible combinations of the three actions. Define $\pi(a_j; \mathbf{s}) = \exp(\mathbf{s}^\top \beta_j) / \{1 + \sum_{k=1}^7 \exp(\mathbf{s}^\top \beta_k)\}$ for $j = 1, \dots, 7$, and $\pi(a_8; \mathbf{s}) = 1 / \{1 + \sum_{k=1}^7 \exp(\mathbf{s}^\top \beta_k)\}$. This defines a class of randomized policies parametrized by $\beta = (\beta_1^\top, \dots, \beta_7^\top)^\top$, where β_j is a vector of parameters for the j -th treatment combination. Actions are selected stochastically according to the probabilities $\pi(a_j; \mathbf{s})$, $j = 1, \dots, 8$. This class of policies is the same as that introduced in Section 4 with $J = 8$. Our implementation of V-learning follows the setup in Section 2 with Ω equal to the identity matrix. Estimation of $\mu^t(a^t; \mathbf{s}^t)$ is done using multinomial logistic regression. Maximizing $\hat{V}_{n, \mathcal{R}}(\pi)$ is done using a combination of simulated annealing and the BFGS algorithm as implemented in the `optim` function in R software. Simulated annealing with 1000 function evaluations is used to find a neighborhood of the maximum; the result is then used as the starting value for the BFGS algorithm. We use this strategy to avoid local maxima.

Although we maximize value over a class of randomized policies, the true optimal policy is deterministic. To prevent the coefficients of $\hat{\beta}_n$ from diverging to infinity, we add an L2 penalty when maximizing over β . To prevent overfitting $V(\pi, \mathbf{s}; \theta^\pi)$, we use an L2 penalty when estimating $\hat{\theta}_n^\pi$, i.e., $\mathcal{P}(\theta^\pi) = \|\theta^\pi\|^2$. Tuning parameters can be used to control the amount of randomness in the estimated policy. For example, increasing the penalty when estimating β is one way to encourage exploration through the action space, because $\beta = 0$ defines a policy where each action is selected with equal probability.

As a comparison, we implement an estimating equation form of infinite horizon Q-learning (Maei et al., 2010; Ertefaie, 2014; Murphy et al., 2016). Define

$$Q^\pi(\mathbf{s}^t, a^t) = \mathbb{E} \left\{ \sum_{k=0}^{\infty} \gamma^k U^{t+k}(\pi) \mid \mathbf{S}^t = \mathbf{s}^t, A^t = a^t \right\}.$$

The Bellman optimality equation is

$$Q^{\text{opt}}(\mathbf{s}^t, a^t) = \mathbb{E} \left\{ U^t + \gamma \max_{a \in \mathcal{A}} Q^{\text{opt}}(\mathbf{S}^{t+1}, a) \mid \mathbf{S}^t = \mathbf{s}^t, A^t = a^t \right\}.$$

Let $Q(\mathbf{s}, a; \eta^{\text{opt}})$ be a parametric model for $Q^{\text{opt}}(\mathbf{s}, a)$ indexed by $\eta^{\text{opt}} \in H \subseteq \mathbb{R}^q$. The Bellman optimality equation motivates the estimating equation

$$0 = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^{T_i} \left\{ u_i^t + \gamma \max_{a \in \mathcal{A}} Q(\mathbf{s}_i^{t+1}, a; \eta^{\text{opt}}) - Q(\mathbf{s}_i^t, a_i^t; \eta^{\text{opt}}) \right\} \nabla_{\eta^{\text{opt}}} Q(\mathbf{s}_i^t, a_i^t; \eta^{\text{opt}}).$$

After solving for the estimate $\hat{\eta}_n^{\text{opt}}$, the optimal policy in state \mathbf{s} is to select action $\hat{\pi}_n(\mathbf{s}) = \arg\max_{a \in \mathcal{A}} Q(\mathbf{s}, a; \hat{\eta}_n^{\text{opt}})$. In our implementation, we model $Q(\mathbf{s}, a; \eta)$ as a linear function with interactions between all state variables and all possible treatments.

Data for each simulated patient are generated as follows. For each day, consisting of 24 time points, we generate 4 time points for food intake, 4 time points for strenuous activity, and 8 time points for moderate activity. Food intake in grams and activity counts are generated for each selected time point and the surrounding time points to reflect the clustering of these activities throughout the day. Data on these lifestyle considerations are generated from normal distributions with means estimated from the data of Maahs et al. (2012). At each time point, each patient is randomly assigned to receive an insulin injection with probability 0.22, estimated from the data of Maahs et al. (2012).

Initial blood glucose level for each patient is drawn from a normal distribution with mean 100 and standard deviation 25. Define the covariates for patient i collected at time t by $(\text{Gl}_i^t, \text{Di}_i^t, \text{Ex}_i^t)^\top$, where Gl_i^t is average blood glucose level, Di_i^t is total dietary intake, and Ex_i^t is total counts of physical activity. Glucose levels evolve according to

$$\text{Gl}^t = \mu(1 - \alpha_1) + \alpha_1 \text{Gl}^{t-1} + \alpha_2 \text{Di}^{t-1} + \alpha_3 \text{Di}^{t-2} + \alpha_4 \text{Ex}^{t-1} + \alpha_5 \text{Ex}^{t-2} + \alpha_6 \text{In}^{t-1} + \alpha_7 \text{In}^{t-2} + e,$$

where In^t is an indicator of an insulin injection received at time t and $e \sim N(0, \sigma^2)$. We use the parameter vector $\alpha = (\alpha_1, \dots, \alpha_7)^\top = (0.9, 0.1, 0.1, -0.005, -0.005, -1, -2)^\top$, $\mu = 100$, and $\sigma = 5.5$ based on a linear model fit to the data of Maahs et al. (2012). The known lag-time in the effect of insulin is reflected by $\alpha_6 = -1$ and $\alpha_7 = -2$. Selecting $\alpha_1 < 1$ ensures an approximately stationary Markov chain. We consider three different models for the state-value function: The first is a linear model, i.e., $\Phi(\mathbf{s}_i^t) = \mathbf{s}_i^t$. We then consider second degree polynomial basis functions and Gaussian radial basis functions (RBF). The Gaussian RBF is $\phi(x; \kappa, \tau^2) = \exp\{-(x - \kappa)^2/2\tau^2\}$. We use $\tau = 0.5$ and $\kappa = 0, 0.5, 1$ to create a basis of functions.

Let n denote number of patients and T denote number of time points per patient. Our choices for n and T are based on what is feasible for an mHealth outpatient study (dietary data was collected on two days by Maahs et al., 2012). For each replication, the optimal treatment regime is estimated using Q-learning and value is estimated using an independent sample of 50 patients with treatments generated according to the estimated optimal regime. We use the parametric value estimate for the V-learning regime, i.e., $\hat{V}_{n, \mathcal{R}}(\hat{\pi}_n) = \hat{\nu}^\top \hat{\theta}_n^{\hat{\pi}_n}$, where $\hat{\nu} = \mathbb{E}_n \Phi(\mathbf{S}^1)$ is the estimated mean of \mathbf{S} under the reference distribution, estimated from the initial states of each simulated data set.

In the setting of type 1 diabetes, we expect delayed effects of treatment to have a short duration due to the volatility of blood glucose levels. Thus, we tend to use smaller discount

factors throughout. Table 2 contains estimated values for V-learning (VL) with different basis functions and Q-learning (QL). Estimated values are averaged over 50 replications. Estimated standard error of the estimated value is in parentheses; Monte Carlo standard errors were calculated as the standard deviation of estimated values across the 50 replications. The simulations in Table 2 were performed using discount factor $\gamma = 0.5$ to place roughly equal weight on short- and long-term utilities. Average value observed in the simulated data is reported alongside estimated value for each estimated treatment regime. We note that

n	T	Linear VL	Polynomial VL	Gaussian VL	QL	Observed
30	24	-0.84 (0.57)	-0.86 (0.53)	-0.89 (0.63)	-1.97 (0.70)	-1.79
	48	-0.65 (0.44)	-0.64 (0.54)	-0.80 (0.54)	-1.88 (0.59)	-1.47
60	24	-1.17 (0.52)	-0.75 (0.75)	-0.89 (0.76)	-1.84 (0.76)	-1.80
	48	-0.66 (0.61)	-0.55 (0.58)	-0.59 (0.60)	-1.90 (0.62)	-1.43

Table 2: Estimated values for V-learning and Q-learning.

V-learning is able to estimate policies with larger estimated value than infinite horizon Q-learning. The increased flexibility conferred by nonlinear basis functions can yield increased estimated values. Estimated value tends to increase with larger n and T and estimated standard error tends to decrease with larger n and T .

Next, we examine how the performance of Q- and V-learning is affected by the choice of discount factor. Table 3 contains estimated values for different γ and T based on 50 replications of the same generative model with $n = 30$. V-learning is implemented with a linear basis. Because of the short duration of delayed effects in this setting, V-learning tends

γ	T	VL	QL	Observed
0.25	24	-1.31 (0.60)	-1.67 (0.71)	-1.80
	48	-1.16 (0.50)	-2.10 (0.69)	-1.45
0.75	24	-2.62 (1.61)	-1.97 (0.78)	-1.79
	48	-2.22 (1.32)	-1.96 (0.62)	-1.45

Table 3: Estimated values for different discount factors.

to perform better using smaller discount factors. Again we see that estimated value tends to increase with larger T and estimated standard error tends to decrease with larger T .

5.2 Accumulating data

In practice, it may be useful for patients with type 1 diabetes to follow a dynamic treatment regime that is updated as new data are collected. Here we consider a hypothetical study wherein n patients are followed for a burn-in period of T' time points, an optimal policy is estimated, and patients are followed for an additional $T - T'$ time points. At each time point, $t \geq T'$, actions are taken according to the most recently estimated policy. Recall that V-learning produces a randomized decision rule from which to sample actions at each time

point. When selecting an action based on a Q-learning policy, we incorporate an ϵ -greedy strategy by selecting the action recommended by the estimated policy with probability $1 - \epsilon$ and otherwise randomly selecting one of the other actions. In our simulations, we use $\epsilon = 0.1$. The discount factor is $\gamma = 0.5$ throughout. The generative model is the same as that above. For the first T' time points, actions are generated randomly from the eight available. We assume throughout that when a patient is recommended a certain action, e.g., food intake, the patient is available and complies. Grams of food intake and activity counts are generated from normal distributions with means estimated from the data of Maahs et al. (2012).

We estimate the first policy after a burn-in period of 10 time points. A new policy is estimated every 5 time points. After T time points, we estimate the value as the average utility over all patients and all time points after the burn-in period. Results in Table 4 are averaged over 50 replications of a hypothetical study. We note that V-learning produces

n	T	linear VL	polynomial VL	Gaussian VL	QL
30	24	-1.18	-1.13	-1.16	-1.91
	48	-1.14	-1.15	-1.11	-2.48
60	24	-1.11	-1.13	-1.13	-2.02
	48	-1.10	-1.12	-1.13	-2.64

Table 4: Estimated values for online V-learning.

larger values than Q-learning when implemented as data accumulate. Again, increased sample size and flexibility in the choice of basis functions result in larger values and estimated value increases with additional time points.

6 Case study: Type 1 diabetes

Machine learning is currently under consideration in type 1 diabetes through studies to build and test a “closed loop” system that joins continuous blood glucose monitoring and subcutaneous insulin infusion through an underlying algorithm. Known as the artificial pancreas, this technology has been shown to be safe in preliminary studies and is making headway from small hospital-based safety studies to large-scale outpatient effectiveness studies (Ly et al., 2014, 2015). Despite the success of the artificial pancreas, the rate of uptake may be limited and widespread use may not occur for many years (Kowalski, 2015). The proposed method may be useful for implementing mHealth interventions for use alongside the artificial pancreas or before it is widely available.

Studies have shown that data on food intake and physical activity to inform optimal decision making can be collected in an inpatient setting (see, e.g., Cobry et al., 2010; Wolever and Mullan, 2011). However, Maahs et al. (2012) demonstrated that rich data on the effect of food intake and physical activity can be collected in an outpatient setting using mobile technology. Here, we apply the proposed methodology to the observational data collected by Maahs et al. (2012).

The full data consist of $N = 31$ patients with type 1 diabetes, aged 12–18. Glucose levels were monitored using continuous glucose monitoring and physical activity tracked using accelerometers for five days. Dietary data were self-reported by the patient in telephone-based interviews for two days. Patients were treated using either an insulin pump or multiple daily insulin injections. We use data on a subset of $n = 14$ patients treated with an insulin pump for whom full follow-up is available on days when dietary information was recorded. This represents 28 patient-days of data, with which we use V-learning to estimate an optimal treatment policy.

The setup closely follows the simulation experiments in Section 5.1. Patient state at each time, t , is taken to be average glucose level and total counts of physical activity over the two previous 60 minutes intervals and total food intake in grams over the four previous 60 minute intervals to allow for delayed effects. The actions available at each time are insulin injection, food intake, exercise, and all combinations of these. The utility at each t is a weighted sum of glycemic events over the 60 minutes preceding and following time t with weights defined in Table 1 in Section 5.1. We seek a treatment regime with large value, i.e., a treatment regime to minimize the number of hypo- and hyperglycemic episodes weighted to reflect the clinical importance of each.

We note that when $V(\pi, \mathbf{s}; \theta^\pi)$ is linear in θ^π , we can evaluate $\hat{V}_{n, \mathcal{R}}(\pi)$ with only the mean of \mathbf{S} under \mathcal{R} . These were estimated from the data. An advantage of V-learning is that an estimate for the value of the optimal policy follows directly from the reference distribution and the estimated parameters.

Table 5 contains estimated values of treatment policies estimated using V-learning with different discount factors and basis functions. Average value observed in the data is -2.59 .

basis	$\gamma = 0.2$	$\gamma = 0.4$
linear	-2.54	-1.89
polynomial	-1.58	-1.88
Gaussian	-1.60	-1.91

Table 5: Values for linear V-learning estimated from data.

The results in Table 5 indicate that improvements in glycemic control can result from dynamic treatment strategies which account for food intake and physical activity; the estimated treatment regimes yield larger values than that observed in the data. Because delayed effects of treatment should have short duration due to the volatility of blood glucose, the best results are achieved when using smaller discount factors. The best treatment regime is estimated using polynomial V-learning with $\gamma = 0.2$. The resulting value is -1.58 ; such an increase in value over that observed in the data could correspond to, e.g., a reduction of two hypoglycemic ($\text{Gl} \leq 70$ mg/dL) events per day.

As an example, consider a patient who presented with an average blood glucose level of 152 mg/dL in the previous hour and 137 mg/dL in hour preceding that. The estimated policy recommends actions according to the probabilities in Table 6. Because this patient presented with moderately high glucose which has been increasing over the past two hours, the policy places highest probabilities on actions that will help to lower blood glucose.

action	probability
no action	0.0008
exercise	0.0024
food	0.0088
food, exercise	0.0050
insulin	0.5005
insulin, exercise	0.4506
insulin, food	0.0195
insulin, food, exercise	0.0122

Table 6: Estimated treatment regime for hyperglycemic patient.

As a second example, consider a patient who presented with an average blood glucose level of 76 mg/dL in the previous hour and 108 mg/dL in the hour preceding that. The estimated policy recommends actions according to the probabilities in Table 7. Because the

action	probability
no action	0.0192
exercise	0.0320
food	0.6435
food, exercise	0.0443
insulin	0.0599
insulin, exercise	0.0695
insulin, food	0.0724
insulin, food, exercise	0.0592

Table 7: Estimated treatment regime for hypoglycemic patient.

patient presented with a borderline hypoglycemic episode, the estimated policy places the most weight on food consumption to prevent the patient’s blood glucose from continuing to drop. Instances of low blood glucose in the observed data are relatively rare. Therefore, there is more uncertainty around the optimal decision for this patient, which is reflected by the more uniform distribution on the action space. Future work in this area may include accounting for insulin on board, the amount of insulin that is active in the patient’s body; this may result in a treatment regime with further improved value.

7 Conclusion

The emergence of mHealth has provided great potential for the estimation and implementation of dynamic treatment regimes in a variety of clinical areas. Mobile technology can be used both in the collection of rich longitudinal data to inform decision making and the delivery of deeply tailored interventions. The proposed method, V-learning, addresses a number of challenges associated with estimating dynamic treatment regimes in mHealth applications.

V-learning directly estimates a policy that maximizes value over a class of policies, requiring minimal assumptions on the data-generating model. Furthermore, V-learning permits estimation of a randomized decision rule which can be used in place of an ϵ -greedy strategy to encourage exploration in online estimation. Estimation of an optimal policy for different populations can be handled through the use of different reference distributions and an estimated value for any reference distribution can be computed directly from the estimated parameters. V-learning and mobile technology have the potential to improve patient outcomes in a variety of clinical areas. We have demonstrated, for example, that the proposed method can be used to reduce the number of hypo- and hyperglycemic episodes in patients with type 1 diabetes. The proposed method is also applicable outside of mHealth applications. For example, V-learning could be used to estimate dynamic treatment regimes from electronic health records data. Future research in this area may include increasing flexibility through use of a semiparametric model for the state-value function. Alternatively, nonlinear models for the state-value function may be informed by underlying theory or mathematical models of the system of interest. Finally, accounting for patient availability and feasibility of a sequence of treatments can be done by setting constraints on the class of policies. This will ensure that the resulting mHealth intervention is able to be implemented and that the recommended decisions are consistent with domain knowledge.

Acknowledgments

We gratefully acknowledge funding sources NIH P01 CA142538, NIH UL1 TR001111, NIH T32 CA201159, NSF DMS-1555141, NSF DMS-1513579, NIH R01 AA023187

References

- Ali, A. A., S. M. Hossain, K. Hovsepian, M. M. Rahman, K. Plarre, and S. Kumar (2012). mpuff: automated detection of cigarette smoking puffs from respiration measurements. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks*, pp. 269–280. ACM.
- Bergenstal, R. M., S. Garg, S. A. Weinzimer, B. A. Buckingham, B. W. Bode, W. V. Tamborlane, and F. R. Kaufman (2016). Safety of a hybrid closed-loop insulin delivery system in patients with type 1 diabetes. *Journal of the American Medical Association* 316(13), 1407–1408.
- Bexelius, C., M. Löf, S. Sandin, Y. T. Lagerros, E. Forsum, and J.-E. Litton (2010). Measures of physical activity using cell phones: validation using criterion methods. *Journal of Medical Internet Research* 12(1), e2.
- Cobry, E., K. McFann, L. Messer, V. Gage, B. VanderWel, L. Horton, and H. P. Chase (2010). Timing of meal insulin boluses to achieve optimal postprandial glycemic control in patients with type 1 diabetes. *Diabetes Technology & Therapeutics* 12(3), 173–177.

- Ertefaie, A. (2014). Constructing dynamic treatment regimes in infinite-horizon settings. *arXiv preprint arXiv:1406.0764*.
- Free, C., G. Phillips, L. Watson, L. Galli, L. Felix, P. Edwards, V. Patel, and A. Haines (2013). The effectiveness of mobile-health technologies to improve health care service delivery processes: a systematic review and meta-analysis. *PLoS Med* 10(1), e1001363.
- Haller, M. J., M. S. Stalvey, and J. H. Silverstein (2004). Predictors of control of diabetes: monitoring may be the key. *The Journal of Pediatrics* 144(5), 660–661.
- Klasnja, P., E. B. Hekler, S. Shiffman, A. Boruvka, D. Almirall, A. Tewari, and S. A. Murphy (2015). Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology* 34(S), 1220.
- Kosorok, M. R. (2008). *Introduction to Empirical Processes and Semiparametric Inference*. New York: Springer.
- Kosorok, M. R. and E. E. Moodie (2015). *Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine*, Volume 21. SIAM.
- Kowalski, A. (2015). Pathway to artificial pancreas systems revisited: moving downstream. *Diabetes Care* 38(6), 1036–1043.
- Kumar, S., W. J. Nilsen, A. Abernethy, A. Atienza, K. Patrick, M. Pavel, W. T. Riley, A. Shar, B. Spring, D. Spruijt-Metz, et al. (2013). Mobile health technology evaluation: the mhealth evidence workshop. *American Journal of Preventive Medicine* 45(2), 228–236.
- Lai, T. L. and H. Robbins (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6(1), 4–22.
- Levine, B.-S., B. J. Anderson, D. A. Butler, J. E. Antisdel, J. Brackett, and L. M. Laffel (2001). Predictors of glycemic control and short-term adverse outcomes in youth with type 1 diabetes. *The Journal of Pediatrics* 139(2), 197–203.
- Liao, P., P. Klasnja, A. Tewari, and S. A. Murphy (2015). Sample size calculations for micro-randomized trials in mhealth. *Statistics in Medicine*.
- Ly, T. T., M. D. Breton, P. Keith-Hynes, D. De Salvo, P. Clinton, K. Benassi, B. Mize, D. Chernavsky, J. Place, D. M. Wilson, et al. (2014). Overnight glucose control with an automated, unified safety system in children and adolescents with type 1 diabetes at diabetes camp. *Diabetes Care* 37(8), 2310–2316.
- Ly, T. T., A. Roy, B. Grosman, J. Shin, A. Campbell, S. Monirabbasi, B. Liang, R. von Eyben, S. Shanmugham, P. Clinton, et al. (2015). Day and night closed-loop control using the integrated medtronic hybrid closed-loop system in type 1 diabetes at diabetes camp. *Diabetes Care* 38(7), 1205–1211.

- Maahs, D. M., E. Mayer-Davis, F. K. Bishop, L. Wang, M. Mangan, and R. G. McMurray (2012). Outpatient assessment of determinants of glucose excursions in adolescents with type 1 diabetes: proof of concept. *Diabetes Technology & Therapeutics* 14(8), 658–664.
- Maei, H. R., C. Szepesvári, S. Bhatnagar, and R. S. Sutton (2010). Toward off-policy learning control with function approximation. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 719–726.
- Moodie, E. E., T. S. Richardson, and D. A. Stephens (2007). Demystifying optimal dynamic treatment regimes. *Biometrics* 63(2), 447–455.
- Murphy, S., Y. Deng, E. Laber, H. Maei, R. Sutton, and K. Witkiewitz (2016). A batch, off-policy, actor-critic algorithm for optimizing the average reward. *arXiv preprint arXiv:1607.05047*.
- Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 65(2), 331–355.
- Murphy, S. A. (2005). A generalization error for q-learning. *Journal of Machine Learning Research* 6(Jul), 1073–1097.
- Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Quinn, C. C., M. D. Shardell, M. L. Terrin, E. A. Barr, S. H. Ballew, and A. L. Gruber-Baldini (2011). Cluster-randomized trial of a mobile phone personalized behavioral intervention for blood glucose control. *Diabetes Care* 34(9), 1934–1942.
- Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium on Biostatistics*, pp. 189–326. Springer.
- Rubin, D. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 34–58.
- Schulte, P., A. Tsiatis, E. Laber, and M. Davidian (2014). Q- and a-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science* (In press), 1–77.
- Steinhubl, S. R., E. D. Muse, and E. J. Topol (2013). Can mobile health technologies transform health care? *Journal of the American Medical Association* 310(22), 2395–2396.
- Sutton, R. and A. Barto (1998). *Reinforcement Learning: An Introduction*. The MIT Press.
- Tang, Y. and M. R. Kosorok (2012). Developing adaptive personalized therapy for cystic fibrosis using reinforcement learning. *The University of North Carolina at Chapel Hill Department of Biostatistics Technical Report Series Working paper 30*.

Weinzimer, S. A., G. M. Steil, K. L. Swan, J. Dziura, N. Kurtz, and W. V. Tamborlane (2008). Fully automated closed-loop insulin delivery versus semiautomated hybrid control in pediatric patients with type 1 diabetes using an artificial pancreas. *Diabetes Care* 31(5), 934–939.

Wolever, T. and Y. Mullan (2011). Sugars and fat have different effects on postprandial glucose responses in normal and type 1 diabetic subjects. *Nutrition, Metabolism and Cardiovascular Diseases* 21(9), 719–725.

Zhao, Y., M. R. Kosorok, and D. Zeng (2009). Reinforcement learning design for cancer clinical trials. *Statistics in Medicine* 28(26), 3294–3315.

Ziegler, R., B. Heidtmann, D. Hilgard, S. Hofer, J. Rosenbauer, and R. Holl (2011). Frequency of smbg correlates with hba1c and acute complications in children and adolescents with type 1 diabetes. *Pediatric Diabetes* 12(1), 11–17.

Appendix

Proof of Lemma 2.1. Let π be an arbitrary policy and $\gamma \in (0, 1)$ a fixed constant. Suppose we observe a state $\mathbf{S}^t = \mathbf{s}^t$ at time t and let $\bar{a}^{t-1} = (a^1, \dots, a^{t-1})$ be the sequence of actions resulting in $\mathbf{S}^t = \mathbf{s}^t$, i.e., $\mathbf{S}^{*t}(\bar{a}^{t-1}) = \mathbf{s}^t$. Let $\bar{a}^{k+1} = (a^t, \dots, a^{t+k}) \in \mathcal{A}^{k+1}$ be a potential sequence of actions taken from time t to time $t+k$. We have that

$$\begin{aligned}
V(\pi, \mathbf{s}^t) &= \sum_{k \geq 0} \gamma^k \mathbb{E} \left\{ U^{*(t+k)}(\pi) \middle| \mathbf{S}^t = \mathbf{s}^t \right\} \\
&= \sum_{k \geq 0} \gamma^k \mathbb{E} \left(\sum_{\bar{a}^{t+k} \in \mathcal{A}^{t+k}} U^{*(t+k)}(\bar{a}^{t+k}) \prod_{v=t}^{t+k} 1 [\xi_{\pi}^v \{ \mathbf{S}^{*v}(\bar{a}^{v-1}) \} = a^v] \middle| \mathbf{S}^t = \mathbf{s}^t \right) \\
&= \sum_{k \geq 0} \gamma^k \sum_{\bar{a}^{k+1} \in \mathcal{A}^{k+1}} U^{*(t+k)}(\bar{a}^{t-1}, \bar{a}^{k+1}) \left\{ \prod_{v=t}^{t+k} \mathbb{E} \left(1 [\xi_{\pi}^v \{ \mathbf{S}^{*v}(\bar{a}^{v-1}) \} = a^v] \middle| \mathbf{S}^t = \mathbf{s}^t \right) \right\} \\
&= \sum_{k \geq 0} \gamma^k \sum_{\bar{a}^{k+1} \in \mathcal{A}^{k+1}} U^{*(t+k)}(\bar{a}^{t-1}, \bar{a}^{k+1}) \prod_{v=t}^{t+k} \pi \{ a^v; \mathbf{S}^{*v}(\bar{a}^{v-1}) \} \prod_{v=t}^{t+k} \frac{\mu^v \{ a^v; \mathbf{S}^{*v}(\bar{a}^{v-1}) \}}{\mu^v \{ a^v; \mathbf{S}^{*v}(\bar{a}^{v-1}) \}} \\
&= \sum_{k \geq 0} \gamma^k \mathbb{E} \left[U^{t+k} \left\{ \prod_{v=0}^k \frac{\pi(a^{t+v}; \mathbf{s}^{t+v})}{\mu^{t+v}(a^{t+v}; \mathbf{s}^{t+v})} \right\} \middle| \mathbf{S}^t = \mathbf{s}^t \right],
\end{aligned}$$

where we let $\pi(a^t; \mathbf{s}^t) = 0$ for all a^t and \mathbf{s}^t whenever $t > T^*(\pi)$. The last equality uses the consistency and strong ignorability assumptions. \square

Proof of Theorem 4.2. Proof of part 1: We first note that θ_0^π must solve

$$0 = \mathbb{E} \left(\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} [U^t + \{\gamma\Phi(\mathbf{S}^{t+1}) - \Phi(\mathbf{S}^t)\}^\top \theta^\pi] \Phi(\mathbf{S}^t) \right),$$

or

$$\mathbb{E} \left[\frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} \Phi(\mathbf{S}^t) \{\Phi(\mathbf{S}^t) - \gamma\Phi(\mathbf{S}^{t+1})\}^\top \right] \theta^\pi = \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} U^t \Phi(\mathbf{S}^t) \right\},$$

which is equivalent to $w_1(\pi)\theta^\pi = w_2(\pi)$ where $w_2(\pi) = \mathbb{E} \{\pi(A^t; \mathbf{S}^t) \mu^t(A^t; \mathbf{S}^t)^{-1} U^t \Phi(\mathbf{S}^t)\}$. We have that

$$\left\| \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} U^t \Phi(\mathbf{S}^t) \right\} \right\| \leq c_0^{-1} (\mathbb{E}|U^t|^2)^{1/2} (\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2)^{1/2} < \infty,$$

by assumption 3, part 1 of assumption 4 and the Cauchy–Schwarz inequality. Let $c \in \mathbb{R}$ be arbitrary and note that

$$\begin{aligned} \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} c^\top \Phi(\mathbf{S}^t) \Phi(\mathbf{S}^{t+1})^\top c \right\} \\ \leq \left[\mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} c^\top \Phi(\mathbf{S}^t)^{\otimes 2} c \right\} \cdot \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} c^\top \Phi(\mathbf{S}^{t+1})^{\otimes 2} c \right\} \right]^{1/2}, \end{aligned}$$

by the Cauchy–Schwarz inequality, where $u^{\otimes 2} = uu^\top$. This implies that

$$\begin{aligned} c^\top w_1(\pi) c &\geq \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} c^\top \Phi(\mathbf{S}^t)^{\otimes 2} c \right\} \\ &\quad - \mathbb{E} \left\{ \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} c^\top \Phi(\mathbf{S}^t)^{\otimes 2} c \right\}^{1/2} \mathbb{E} \left\{ \gamma^2 \frac{\pi(A^t; \mathbf{S}^t)}{\mu^t(A^t; \mathbf{S}^t)} c^\top \Phi(\mathbf{S}^{t+1})^{\otimes 2} c \right\}^{1/2} \\ &= A - A^{1/2} B^{1/2} \\ &= A^{1/2} (A^{1/2} - B^{1/2}) \\ &= \frac{A^{1/2} (A - B)}{A^{1/2} + B^{1/2}}, \end{aligned}$$

where we simplify notation by defining $A = \mathbb{E} \{\pi(A^t; \mathbf{S}^t) \mu^t(A^t; \mathbf{S}^t)^{-1} c^\top \Phi(\mathbf{S}^t)^{\otimes 2} c\}$ and $B = \mathbb{E} \{\gamma^2 \pi(A^t; \mathbf{S}^t) \mu^t(A^t; \mathbf{S}^t)^{-1} c^\top \Phi(\mathbf{S}^{t+1})^{\otimes 2} c\}$. We have that

$$\begin{aligned} A^{1/2} + B^{1/2} &\leq c_0^{-1/2} \|c\| \{\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2\}^{1/2} + c_0^{-1/2} \|c\| \{\mathbb{E}\|\Phi(\mathbf{S}^{t+1})\|^2\}^{1/2} \\ &= 2c_0^{-1/2} \|c\| \{\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2\}^{1/2} \\ &< \infty, \end{aligned}$$

where the equality follows because $\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2 = \mathbb{E}\|\Phi(\mathbf{S}^{t+1})\|^2$ by time-homogeneity. Also, $A \geq A - B$ and $A - B \geq c_1\|c\|^2$ by assumption 5. Thus,

$$\begin{aligned} A - A^{1/2}B^{1/2} &\geq \frac{c_1^{3/2}\|c\|^3}{2c_0^{-1/2}\|c\|\{\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2\}^{1/2}} \\ &= \frac{c_0^{1/2}c_1^{3/2}\|c\|^2}{2\{\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2\}^{1/2}}, \end{aligned}$$

which finally implies that $\theta_0^\pi = w_1(\pi)^{-1}w_2(\pi)$ is well-defined uniformly over $\pi \in \Pi$. Using the fact that $c^\top w_1(\pi)c \geq k_0\|c\|^2$ for a constant $k_0 > 0$, we can show that $\|w_1(\pi)^{-1}\| \leq k_1^{-1}$ for some constant $k_1 > 0$, where $\|\cdot\|$ is the usual matrix norm when applied to a matrix. Therefore, $\|\theta_0^\pi\| \leq k_1^{-1}\|w_2(\pi)\| \leq c_0^{-1}k_1^{-1}\{\mathbb{E}(U^t)^2\}^{1/2}\{\mathbb{E}\|\Phi(\mathbf{S}^t)\|^2\}^{1/2} < \infty$. Finally, it follows from assumptions 5 and 7 that $\sup_{\|\beta_1 - \beta_2\| \leq \delta} \|\theta^{\pi_{\beta_1}} - \theta^{\pi_{\beta_2}}\| \rightarrow 0$ as $\delta \downarrow 0$.

Proof of part 2: Define

$$\mathcal{G} = \left\{ \Phi(\mathbf{s}^t)\Phi(\mathbf{s}^t)^\top / \mu^t(a^t; \mathbf{s}^t), \gamma\Phi(\mathbf{s}^t)\Phi(\mathbf{s}^{t+1})^\top / \mu^t(a^t; \mathbf{s}^t), u^t\Phi(\mathbf{s}^t) / \mu^t(a^t; \mathbf{s}^t) \right\}.$$

Let G be an envelope for \mathcal{G} , for example $G(\mathbf{s}^{t+1}, a^t, \mathbf{s}^t) = \max_{g \in \mathcal{G}} g(\mathbf{s}^{t+1}, a^t, \mathbf{s}^t)$. By part 1 of assumption 4, $\|G\|_{3\rho/2} < \infty$. Part 4 of Theorem 7.1 below gives us that \mathcal{G} is Donsker. Since Π satisfies $J_\Pi\{\infty, \Pi, L_{3\rho}(P)\} < \infty$, we have that

$$\mathcal{F}_1 = \left\{ \frac{\pi(a^t; \mathbf{s}^t)}{\mu^t(a^t; \mathbf{s}^t)} \Phi(\mathbf{s}^t) \{\Phi(\mathbf{s}^t) - \gamma\Phi(\mathbf{s}^{t+1})\}^\top : \pi \in \Pi \right\}$$

satisfies $J_\Pi\{\infty, \mathcal{F}_1, L_{3\rho}(P)\} < \infty$ by part 1 of Lemma 7.1 below. Moreover, $F(a^t, \mathbf{s}^t) = \|\Phi(\mathbf{s}^t)\| \cdot \|\Phi(\mathbf{s}^t) - \gamma\Phi(\mathbf{s}^{t+1})\| / \mu^t(a^t; \mathbf{s}^t)$ is an envelope for \mathcal{F}_1 with $\mathbb{E}F^{3\rho} < \infty$ by part 1 of assumption 4 and assumption 3. Thus, \mathcal{F}_1 is Donsker. Let

$$\mathcal{F}_2 = \left\{ \frac{\pi(a^t; \mathbf{s}^t)}{\mu^t(a^t; \mathbf{s}^t)} u^t\Phi(\mathbf{s}^t) : \pi \in \Pi \right\}.$$

Similar arguments yield that \mathcal{F}_2 is Donsker. Now, let $\hat{A}(\pi) = \{\mathbb{E}_n f_{1\pi} : f_{1\pi} \in \mathcal{F}_1\}$ and $\hat{B}(\pi) = \{\mathbb{E}_n f_{2\pi} : f_{2\pi} \in \mathcal{F}_2\}$. Then,

$$\begin{aligned} \sqrt{n}(\hat{\theta}_n^\pi - \theta_0^\pi) &= \sqrt{n} \left\{ \hat{A}(\pi)^{-1} \hat{B}(\pi) - \hat{A}(\pi)^{-1} \hat{A}(\pi) \theta_0^\pi \right\} + o_P(1) \\ &= \hat{A}(\pi)^{-1} \sqrt{n} \left\{ \hat{B}(\pi) - \hat{A}(\pi) \theta_0^\pi \right\} + o_P(1), \end{aligned}$$

where $o_P(1)$ doesn't depend on π . Using similar arguments, one can show that $\mathcal{F}_3 = \{f_{2\pi} - f_{1\pi}\theta : f_{1\pi} \in \mathcal{F}_1, f_{2\pi} \in \mathcal{F}_2, \pi \in \Pi, \theta \in B_*\}$ is Donsker, where B_* is any finite collection of elements of \mathbb{R}^q . By part 1 of this theorem, there exists a bounded, closed set B_0 such that $\hat{\theta}_0^\pi \in B_0$ for all $\pi \in \Pi$. Let $\mathbb{G}_n(\pi, \theta) = \sqrt{n}(\mathbb{E}_n - \mathbb{E})(f_{2\pi} - f_{1\pi}\theta)$. Note that

$$\begin{aligned} \sup_{\pi \in \Pi} \|\mathbb{G}_n(\pi, \theta_1) - \mathbb{G}_n(\pi, \theta_2)\| &\leq \sup_{\pi \in \Pi} \|\sqrt{n}(\mathbb{E}_n - \mathbb{E})f_{1\pi}\| \cdot \|\theta_1 - \theta_2\| \\ &\leq R^* \|\theta_1 - \theta_2\|, \end{aligned}$$

where $R^* = O_P(1)$ by the Donsker property of \mathcal{F}_1 and R^* doesn't depend on π . Thus, $\mathbb{G}_n(\pi, \theta)$ is stochastically equicontinuous on B_0 . Combined with the Donsker property of \mathcal{F}_3 for arbitrary B_* , we have that the class $\mathcal{F}_4 = \{f_{2\pi} - f_{1\pi}\theta : f_{1\pi} \in \mathcal{F}_1, f_{2\pi} \in \mathcal{F}_2, \pi \in \Pi, \theta \in B_0\}$ is Donsker. Using Slutsky's Theorem, Theorem 11.24 of Kosorok (2008) and the fact that \mathcal{F}_1 is Glivenko–Cantelli, we have that $\sqrt{n}(\hat{\theta}_n^\pi - \theta_0^\pi) = \hat{A}(\pi)^{-1}\mathbb{G}_n(\pi, \theta_0^\pi) \rightsquigarrow w_1(\pi)^{-1}\mathbb{G}_1(\pi)$ in $\ell^\infty(\Pi)$, where $\mathbb{G}_1(\pi)$ is a mean zero Gaussian process indexed by Π with covariance $\mathbb{E}\{\mathbb{G}_1(\pi_1)\mathbb{G}_1(\pi_2)\} = w_0(\pi_1, \pi_2)$.

Proof of part 3: We have that

$$\begin{aligned}\sqrt{n}\left\{\hat{V}_{n,\hat{\mathcal{R}}}(\pi) - V_{\mathcal{R}}(\pi)\right\} &= \sqrt{n}\mathbb{E}_n\Phi(\mathbf{S}^t)\left(\hat{\theta}_n^\pi - \theta_0^\pi\right) \\ &\rightsquigarrow \nu^\top w_1(\pi)^{-1}\mathbb{G}(\pi)\end{aligned}$$

in $\ell^\infty(\Pi)$ by Slutsky's Theorem. \square

Proof of Theorem 4.3. Proof of part 1: Following part 3 of Theorem 4.2, we have that $\sup_{\beta \in \mathcal{B}} |\hat{V}_{n,\hat{\mathcal{R}}}(\pi_\beta) - V_{\mathcal{R}}(\pi_\beta)| \xrightarrow{P} 0$. The proof of part 2 now follows from the unique and well separated maximum condition, i.e., assumption 6. The proof of part 3 follows standard arguments. \square

Lemma 7.1. *Let \mathcal{F} and \mathcal{G} be function classes with respective envelopes F and G . Let $\|F\|_u = (\mathbb{E}\|F\|^u)^{1/u}$. For any $1 \leq r, s_1, s_2 \leq \infty$ with $s_1^{-1} + s_2^{-1} = 1$,*

1. $J_\square\{\infty, \mathcal{F} \cdot \mathcal{G}, L_r(P)\} \leq 2(\|F\|_{rs_1} + \|G\|_{rs_2}) [J_\square\{\infty, \mathcal{F}, L_{rs_1}(P)\} + J_\square\{\infty, \mathcal{G}, L_{rs_2}(P)\}]$.
2. $J_\square\{\infty, \mathcal{F} + \mathcal{G}, L_r(P)\} \leq 2 [J_\square\{\infty, \mathcal{F}, L_r(P)\} + J_\square\{\infty, \mathcal{G}, L_r(P)\}]$.
3. *For any $0 \leq r \leq \infty$, $J_\square\{\infty, \mathcal{F} \cup \mathcal{G}, L_r(P)\} \leq \sqrt{\log 2}(\|F\|_r + \|G\|_r) + J_\square\{\infty, \mathcal{F}, L_r(P)\} + J_\square\{\infty, \mathcal{G}, L_r(P)\}$.*
4. *If \mathcal{G} is a finite class, $J_\square\{\infty, \mathcal{G}, L_r(P)\} \leq 2\|G\|_r\sqrt{\log |\mathcal{G}|}$, where $|\mathcal{G}|$ denotes the cardinality of \mathcal{G} .*

Proof of Lemma 7.1. Proof of part 1: Let $1 \leq r, s_1, s_2 \leq \infty$ with $s_1^{-1} + s_2^{-1} = 2$ and let (ℓ_F, u_F) and (ℓ_G, u_G) be $L_{rs_1}(P)$ and $L_{rs_2}(P)$ ϵ -brackets, respectively. Choose $\ell_F \leq f_1, f_2 \leq u_F$ and $\ell_G \leq g_1, g_2 \leq u_G$ and consider the bracket for any f_2g_2 defined by $f_1g_1 \pm (F|u_G - \ell_G| + G|u_F - \ell_F|)$. Note that

$$f_1g_1 + F|u_G - \ell_G| + G|u_F - \ell_F| - f_2g_2 \geq F|u_G - \ell_G| + G|u_F - \ell_F| - F|g_1 - g_2| - G|f_1 - f_2| \geq 0,$$

because $f_2g_2 - f_1g_1 = f_2g_2 + f_2g_1 - f_2g_1 - f_1g_1 \leq F|g_1 - g_2| + G|f_1 - f_2|$. Similarly, $f_2g_2 + F|u_G - \ell_G| + G|u_F - \ell_F| - f_1g_1 \geq 0$. Thus, these brackets hold all f_2g_2 for $f_2 \in (\ell_F, u_F)$ and $g_2 \in (\ell_G, u_G)$. Now, $\|F|u_G - \ell_G| + G|u_F - \ell_F|\|_r \leq \|F\|_{rs_1}\epsilon + \|G\|_{rs_2}\epsilon$ by Minkowski's inequality and Hölder's inequality and it follows that

$$N_\square\{2\epsilon(\|F\|_{rs_1} + \|G\|_{rs_2}), \mathcal{F} \cdot \mathcal{G}, L_r(P)\} \leq N_\square\{\epsilon, \mathcal{F}, L_{rs_1}(P)\} N_\square\{\epsilon, \mathcal{G}, L_{rs_2}(P)\}.$$

Next we note that

$$N_{[]} \{ \epsilon, \mathcal{F} \cdot \mathcal{G}, L_r(P) \} \leq N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs_1} + \|G\|_{rs_2})}, \mathcal{F}, L_{rs_1}(P) \right\} N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs_1} + \|G\|_{rs_2})}, \mathcal{G}, L_{rs_2}(P) \right\}$$

and thus

$$\begin{aligned} J_{[]} \{ \infty, \mathcal{F} \cdot \mathcal{G}, L_r(P) \} &\leq \int_0^{2\|F\|_{rs_1}\|G\|_{rs_2}} \sqrt{\log N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs_1} + \|G\|_{rs_2})}, \mathcal{F}, L_{rs_1}(P) \right\}} d\epsilon \\ &\quad + \int_0^{2\|F\|_{rs_1}\|G\|_{rs_2}} \sqrt{\log N_{[]} \left\{ \frac{\epsilon}{2(\|F\|_{rs_1} + \|G\|_{rs_2})}, \mathcal{G}, L_{rs_2}(P) \right\}} d\epsilon \\ &\leq 2(\|F\|_{rs_1} + \|G\|_{rs_2}) [J_{[]} \{ \infty, \mathcal{F}, L_{rs_1}(P) \} + J_{[]} \{ \infty, \mathcal{G}, L_{rs_2}(P) \}]. \end{aligned}$$

The proof of part 2 follows from Lemma 9.25 part (i) of Kosorok (2008) after a change of variables. Proof of part 3: First note that

$$N_{[]} \{ \epsilon, \mathcal{F} \cup \mathcal{G}, L_r(P) \} \leq N_{[]} \{ \epsilon, \mathcal{F}, L_r(P) \} + N_{[]} \{ \epsilon, \mathcal{G}, L_r(P) \},$$

whence it follows that

$$\begin{aligned} J_{[]} \{ \infty, \mathcal{F} \cup \mathcal{G}, L_r(P) \} &= \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log N_{[]} \{ \epsilon, \mathcal{F} \cup \mathcal{G}, L_r(P) \}} d\epsilon \\ &\leq \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log [N_{[]} \{ \epsilon, \mathcal{F}, L_r(P) \} + N_{[]} \{ \epsilon, \mathcal{G}, L_r(P) \}]} d\epsilon \\ &\leq \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log 2 + \log N_{[]} \{ \epsilon, \mathcal{F}, L_r(P) \} + \log N_{[]} \{ \epsilon, \mathcal{G}, L_r(P) \}} d\epsilon \\ &\leq \int_0^{2(\|F\|_r + \|G\|_r)} \sqrt{\log 2} d\epsilon + J_{[]} \{ \infty, \mathcal{F}, L_r(P) \} + J_{[]} \{ \infty, \mathcal{G}, L_r(P) \}, \end{aligned}$$

where the second inequality uses the fact that $a + b \leq 2ab$ for all $a, b \geq 1$.

Proof of part 4: If \mathcal{G} is finite, then $N_{[]} \{ \epsilon, \mathcal{G}, L_r(P) \} \leq |\mathcal{G}|$. Thus,

$$\begin{aligned} J_{[]} \{ \infty, \mathcal{G}, L_r(P) \} &= \int_0^{2\|G\|_r} \sqrt{\log N_{[]} \{ \epsilon, \mathcal{G}, L_r(P) \}} d\epsilon \\ &\leq \int_0^{2\|G\|_r} \sqrt{\log |\mathcal{G}|} d\epsilon, \end{aligned}$$

which completes the proof. \square

Lemma 7.2. *Define the class of functions*

$$\Pi = \left\{ \pi_{\tilde{\beta}}(a; \mathbf{s}) = \frac{a_J + \sum_{j=1}^{J-1} a_j \exp(\mathbf{s}_j^\top \beta_j)}{1 + \sum_{j=1}^{J-1} \exp(\mathbf{s}_j^\top \beta_j)} : \tilde{\beta} = (\beta_1^\top, \dots, \beta_{J-1}^\top)^\top, \tilde{\beta} \in \mathcal{B} \subset \mathbb{R}^{p(J-1)} \right\}$$

for a compact set \mathcal{B} and $2 \leq J < \infty$ where $a = (a_1, \dots, a_J)^\top$. Then, there exists a $b_0 < \infty$ such that for any $1 \leq r \leq \infty$, $J_{[]} \{ \infty, \Pi, L_r(P) \} \leq b_0 \|\mathbf{S}\|_r \sqrt{p(J-1)\pi}$, which is finite whenever $\|\mathbf{S}\|_r < \infty$. Furthermore, $\sup_{\|\tilde{\beta}_1 - \tilde{\beta}_2\| \leq \delta} \mathbb{E} \|\pi_{\tilde{\beta}_1}(A; \mathbf{S}) - \pi_{\tilde{\beta}_2}(A; \mathbf{S})\| \rightarrow 0$ as $\delta \downarrow 0$.

Proof of Lemma 7.2. For $\tilde{\beta}_1, \tilde{\beta}_2 \in \mathcal{B}$, define $d(\tilde{\beta}_1, \tilde{\beta}_2) = \max_{1 \leq j \leq J-1} \|\tilde{\beta}_{1j} - \tilde{\beta}_{2j}\|$ and $b_0 = \sup_{\tilde{\beta}_1, \tilde{\beta}_2 \in \mathcal{B}} \|\tilde{\beta}_1 - \tilde{\beta}_2\| < \infty$ because \mathcal{B} is compact. By the mean value theorem, for any $\tilde{\beta}_1, \tilde{\beta}_2 \in \mathcal{B}$, there exists a point $\tilde{\beta}_*$ on the line segment between $\tilde{\beta}_1$ and $\tilde{\beta}_2$ such that

$$\pi_{\tilde{\beta}_1}(a; \mathbf{s}) - \pi_{\tilde{\beta}_2}(a; \mathbf{s}) = \frac{1}{1 + \sum_{j=1}^{J-1} \exp(\mathbf{s}^\top \tilde{\beta}_{*j})} \left[\sum_{j=1}^{J-1} \{a_j - \pi_{\tilde{\beta}_*}(a; \mathbf{s})\} \exp(\mathbf{s}^\top \tilde{\beta}_{*j}) \mathbf{s}^\top (\tilde{\beta}_{1j} - \tilde{\beta}_{2j}) \right],$$

which implies that

$$|\pi_{\tilde{\beta}_1}(a; \mathbf{s}) - \pi_{\tilde{\beta}_2}(a; \mathbf{s})| \leq \|\mathbf{s}\| d(\tilde{\beta}_1, \tilde{\beta}_2). \quad (6)$$

It follows from equation (6) that assumption 7 holds for this particular class of policies. Now, $N_{[]} \{2\epsilon \|\mathbf{S}\|_r, \Pi, L_r(P)\} \leq N(\epsilon, \mathcal{B}, d)$ by Theorem 9.23 of Kosorok (2008). Furthermore, $N(\epsilon, \mathcal{B}, d) \leq \max \{(b_0/\epsilon)^{p(J-1)}, 1\}$, and thus

$$\begin{aligned} J_{[]} \{\epsilon, \Pi, L_r(P)\} &\leq 2\|\mathbf{S}\|_r \int_0^{b_0} \sqrt{p(J-1) \{\log b_0 + \log(1/\epsilon)\}} d\epsilon \\ &\leq 2\|\mathbf{S}\|_r b_0 \sqrt{p(J-1)} \int_0^1 \sqrt{\log(1/\epsilon)} d\epsilon \\ &= 2\|\mathbf{S}\|_r b_0 \sqrt{p(J-1)} \int_0^\infty u^{1/2} \exp(-u) du \\ &= \|\mathbf{S}\|_r b_0 \sqrt{p(J-1)} \pi, \end{aligned}$$

which proves the result. \square